

What is claimed is:

1. In an adaptive load balancer, a method for processing client file access transactions, the method comprising the steps of:

- 5 presenting a virtual file system to the client computer system, the virtual file system providing access to an aggregated set of files stored by a plurality of server file systems respectively associated with a plurality of server computer systems;
- receiving a client file access transaction from a client computer system, the client file access transaction specifying a file access operation to be performed relative to the
- 10 virtual file system presented to the client computer system;
- processing the client file access transaction in relation to metadata associated with the virtual file system to perform at least one of:
- if the file access operation requires access to a file stored by the plurality of server file systems, then translating, using the metadata associated with the
- 15 virtual file system, the client file access transaction into a server file access transaction for access to the file maintained within one of the plurality of server computer systems according to the file access operation specified by the client file access transaction; and
- if the file access operation does not require access to a file stored by the
- 20 plurality of server file systems, then accessing metadata associated with the virtual file system to complete processing of the file access operation specified by the client file access transaction.

2. The method of claim 1 wherein if the file access operation requires access to a file

25 stored by the plurality of server file systems, the method comprises:

 forwarding the server file access transaction to a selected one of the plurality of server computer systems specified by the metadata for processing of the file access operation specified by the client file access transaction by the selected one of the plurality of server computer systems;

- 30 receiving a server transaction response from the server computer system to which the server file access transaction was forwarded;

translating, using the metadata associated with the virtual file system, the server transaction response into a client file access response; and

forwarding the client file access response to the client computer system from which the client file access transaction was received.

5

3. The method of claim 2 comprising:

maintaining the metadata associated with the virtual file system in at least one forwarding table containing forwarding table entries that provide a mapping of virtual file system parameters to physical file system parameters, the mapping including:

10 a mapping of virtual file identifiers corresponding to each available file accessible by client computer systems in the virtual file system to:

i) at least one corresponding physical file identifier; and

ii) at least one server computer system of the plurality of server computer systems at which an instance of that file can be accessed using that corresponding physical file

15 identifier.

4. The method of claim 3 wherein translating, using the metadata associated with the virtual file system, the client file access transaction into a server file access transaction comprises:

20 obtaining a virtual file identifier specified in the client file access transaction;

matching the virtual file identifier specified in the client file access transaction to a matching forwarding table entry in the at least one forwarding table to identify a corresponding physical file identifier contained in the matching forwarding table entry;

25 creating the server file access transaction by replacing the virtual file identifier specified in the client file access transaction with the corresponding physical file identifier contained in matching forwarding table entry;

and wherein forwarding the server file access transaction to a selected one of the plurality of server computer systems comprises:

30 identifying, from the matching forwarding table entry, the at least one

server computer system at which an instance of that file can be accessed using the corresponding physical file identifier;

forwarding the server file access transaction to a server computer system corresponding to the identified at least one server computer system at which an instance of that file can be accessed.

- 5 5. The method of claim 4 comprising maintaining an active transaction table that associates client file access transactions to corresponding server file access transactions; and

wherein receiving a client file access transaction comprises:

- 10 assigning a client transaction identity to the received client file access transaction; and
- determining if the active transaction table contains an active transaction table entry that corresponds to the client transaction identity, and if not, creating an active transaction table entry containing the assigned client transaction identity and an identity of the client computer
- 15 system from which the client file access transaction was received;

wherein creating the server file access transaction comprises:

- assigning a server transaction identity to the created server file access transaction; and
- 20 storing the server transaction identity in the active transaction table entry that contains the client transaction identity of the received client file access transaction that was translated to the created server file access transaction.

- 25 6. The method of claim 5 wherein translating, using the metadata associated with the virtual file system, the server transaction response into a client file access response comprises:

- obtaining a physical file identifier specified in the server transaction response;
- matching the physical file identifier specified in the server transaction response to a matching forwarding table entry in the at least one forwarding table to identify a
- 30 corresponding virtual file identifier contained in the matching forwarding table entry;

replacing the physical file identifier specified in the server transaction response with the corresponding virtual file identifier contained in matching forwarding table entry to create the client file access response; and

wherein forwarding the client file access response to the client computer system
5 comprises:

obtaining a server transaction identity specified in the server transaction response;

matching the server transaction identity to an active transaction table entry containing that server transaction identity in the an active transaction table to
10 obtain a client transaction identity that corresponds to that server transaction identity; and

forwarding the client file access response to a client computer system corresponding to a client computer system associated with the client transaction identity identified in matching active transaction table entry in the active
15 transaction table.

7. The method of claim 3 wherein the at least one forwarding table includes:

a directory location mapping table containing a mapping of virtual file identities of parent and child directories within the virtual file system, the parent and child
20 directories representing an aggregation of directory structures present within respective server file systems of the plurality of file servers; and

a set of directory file mapping tables, one for each directory in the virtual file system, each directory file mapping table containing a mapping of virtual file identities of files within the virtual file system of that directory to physical file identities of a
25 corresponding file and server computer system and within an aggregation of the respective server file systems of each of the plurality of file servers.

8. The method of claim 7 wherein processing the client file access transaction in relation to metadata associated with the virtual file system comprises:

30 obtaining a virtual file identity within the client file access transaction;

comparing a directory portion of the virtual file identity to the directory location mapping table to identify a destination virtual file identity of a directory in the virtual file

system that contains a file to be accessed according to the file access operation specified in the client file access transaction;

identifying a specific directory file mapping table within the set of directory file mapping tables that corresponds to the destination virtual file identity of the directory in the virtual file system that contains a file to be accessed according to the file access operation specified in the client file access transaction;

within the specific identified directory file mapping table, comparing a file portion of the virtual file identity to directory file mapping table entries to identify at least one matching directory file mapping table entry that contains the file portion of the virtual file identity in order to identify a server computer system of the plurality of server computer systems and a physical file identity of a file on that server computer system that corresponds to the file to be accessed according to the file access operation specified in the client file access transaction.

9. The method of claim 8 wherein comparing a file portion of the virtual file identity to directory file mapping table entries to identify at least one matching directory file mapping table entry comprises:

identifying a plurality of matching directory file mapping table entries that each contains the file portion of the virtual file identity, thus indicating that the file to be accessed according to the file access operation specified in the client file access transaction is replicated in at least two locations within the plurality of server computer systems.

10. The method of claim 2 wherein the plurality of server computer systems includes:

a first server computer system that provides a first server file system having a first set of files available within the virtual file system;

a second server computer system that provides a second server file system having a second set of files available within the virtual file system; and

wherein:

the first set of files contains different files than the second set of files;

the adaptive load balancer presents the first and second set of files as the aggregated set of files without identifying to the client computer system that the first and second set of files are stored on respective first and second server computer systems; and

the adaptive load balancer uses the metadata associated with the virtual file
5 system to manage access to the first and second set of files within the first and second server computer systems on behalf of client computer systems requesting access to the set of files.

11. The method of claim 10 wherein the virtual file system provided by the adaptive load
10 balancer provides at least one file accessible to the client computer systems that is replicated within the first and second file systems on the first and second server computer systems.

12. The method of claim 2 comprising:

15 detecting a trend of access to a file within the virtual file system by tracking file access patterns to the file by at least one client computer system;

in response to detecting the trend of access, performing a migration operation on the file for which the trend is detected in order to manage access to the file.

20 13. The method of claim 12 wherein detecting a trend of access to a file comprises:

detecting frequent access to a file over a time period;

and wherein performing a migration operation on the file comprises:

replicating the file to at least one server computer system that does not
contain the file within its file system at the time of detecting frequent access to a
25 file.

14. The method of claim 13 wherein replicating the file to at least one server computer system that does not contain the file within its file system at the time of detecting frequent access to a file comprises:

30 selecting, as the at least one server computer system for replicating the file, at least one server computer system that is of higher performance in relation to at least one

other server computer system in the plurality of server computer systems available to store the replicate file.

15. The method of claim 12 wherein detecting a trend of access to a file comprises:

5 detecting infrequent access to a file over a time period; and

wherein performing a migration operation on the file comprises:

 identifying, using metadata associated with the virtual file system, that the file is replicated across a plurality of server file systems of respective server computer systems;

10 in response to identifying that the file is replicated, instructing at least one server computer system that maintains a replicated copy of the file within its server file system to remove the replicated copy.

16. The method of claim 15 wherein instructing at least one server computer system that

15 maintains a replicated copy of the file within its server file system to remove the replicated copy comprises:

 selecting, as the at least one server computer system from which the replicated copy of the file is to be removed, at least one server computer system that contains a replicated copy of the file and that is of higher performance relative to other server computer systems in the plurality of server computer systems.

20

17. The method of claim 12 wherein detecting a trend of access to a file comprises:

 detecting infrequent access to a file over a time period;

and wherein performing a migration operation on the file comprises:

25 migrating the file from a higher performing server computer system to a lower performing server computer system.

18. The method of claim 5 wherein receiving a client file access transaction comprises:

30 identifying that the client file access transaction received from the client computer system is a duplicate client file access transaction within the active transaction table;

in response to identifying that the client file access transaction is a duplicate, ignoring the duplicate file access transaction and bypassing processing the client file access transaction in relation to metadata associated with the virtual file system.

5 19. The method of claim 1 wherein the file access operation specified by the client file access transaction is a metadata-only operation that the adaptive load balancer can perform by accessing the metadata associated with the virtual file system without requiring access to files stored within respective file systems of the plurality of computer systems.

10

20. The method of claim 19 wherein:

the file access operation in the client file access transaction specifies an operation to be performed on directories and files of the virtual file systems that the adaptive load balancer provides access to using the Network File System protocol; and

15

wherein the file access operation is a Network File System operation including at least one of a CREATE operation, a MKDIR operation, a REaddir operation and a STATFS operation.

20

21. The method of claim 20 wherein the file access operation is a STATFS operation and wherein accessing metadata associated with the virtual file system to complete processing of the file access operation specified by the client file access transaction comprises:

calculating an amount of storage space associated with each server file system of the plurality of server computer systems that is provided within the virtual file system; and

25

providing a client file access response to the client computer system that indicates an amount of file storage space available for storage of data within the virtual file system that is based on the calculated amount of storage space associated with each server file system.

30

22. The method of claim 21 wherein calculating an amount of storage space associated with each server file system comprises:

calculating a total amount of current free space within each server file system of the plurality of server computer systems that is provided within the virtual file system;

identifying a replicated amount of storage of replicated files within the set of server file systems maintained by the respective server computer systems within the virtual file system;

calculating a potential amount of total storage space as the total amount of current free space in addition to the replicated amount of storage of replicated files within the set of server file systems; and

setting the amount of file storage space available for storage of data within the virtual file system equal to the potential amount of total storage space.

23. The method of claim 2 wherein the file access operation specified by the client file access transaction is a file access operation that requires modification of a file within the virtual file system; and

wherein receiving a server transaction response from the server computer system to which the server file access transaction was forwarded comprises:

identifying that the server transaction response indicates completion of modification of a file within a first server file system of a first server computer system as specified in the file access operation specified by the client file access transaction;

identifying if replicated copies of the file that was modified exist within server file systems of server computer systems other than the first server computer system;

and if replicated copies of the file that was modified exist, then performing an update replica operation comprising:

i) indicating, within the metadata of the virtual file system associated with the file that was modified, that the file in the first server file system on the first server computer system is a master copy of the file;

ii) indicating, within the metadata of the virtual file system associated with the replicated copies of the file, that each replicated copy is a non-master copy of the file.

24. The method of claim 23 wherein performing the update replica operation further comprises:

incrementally updating each replicated copy of the file that is indicated as a non-master copy of the file such that the replicated copies of the file contain the same content as the file that was modified.

5 25. The method of claim 2 wherein the file access operation specified by the client file access transaction is a file access operation that requires modification of a file within the virtual file system; and

wherein receiving a server transaction response from the server computer system to which the server file access transaction was forwarded comprises:

10 identifying that the server transaction response indicates completion of modification of a file within a first server file system of a first server computer system as specified in the file access operation specified by the client file access transaction;

identifying if replicated copies of the file that was modified exist within server file systems of server computer systems other than the first server computer system; and

15 if replicated copies of the file that was modified exist, then deleting each replicated copy of the file within respective server file systems of server computer systems that maintain a replicated copy of that file, such that file that was modified in the first server file system is the only copy of that file that remains accessible to client computer systems within the virtual file system.

20

26. The method of claim 2 wherein the file access operation specified by the client file access transaction is a file access operation that requires modification of a file within the virtual file system; and

25 wherein translating, using the metadata associated with the virtual file system, the client file access transaction into a server file access transaction for access to the file maintained within one of the plurality of server computer systems comprises:

identifying if replicated copies of the file to be modified exist in the virtual file system within a plurality of server file systems of the plurality of server computer systems; and

30 if replicated copies of the file to be modified exist,

translating, using the metadata associated with the virtual file system, the client file access transaction into a server file access transaction for access to each

copy of the replicated file maintained within the server file systems of the plurality of server computer systems; and wherein forwarding the server file access transaction to a selected one of the plurality of server computer systems specified by the metadata for processing of the file access operation comprises:

5 forwarding the server file access transaction to each of the plurality of server computer systems specified by the metadata that contains a replicated copy of the file to be modified for processing of the file access operation on each replicated copy of the file.

10

27. The method of claim 1 wherein presenting the virtual file system to the client computer system comprises:

aggregating a set of files and directories accessible in the respective file systems of each of the plurality of server computer systems to appear to client computer systems as the virtual file system;

15

exposing the set of files and directories aggregated in the virtual file system to client computer systems using a file access protocol compatible with the client computer systems such that virtual file system is presented to client computer systems as a single virtual file system providing a single point of access to the files and directories of the aggregated set of file and directories of each of the plurality of server computer systems.

20

28. The method of claim 27 wherein the file access protocol in at least one of the Network File System (NFS) protocol and the Common Internet File System (CIFS) protocol and the Hypertext Transfer Protocol and the Simple Object Access Protocol.

25

29. In an adaptive load balancer, a method for processing client data access transactions, the method comprising the steps of:

presenting a virtual data system to the client computer system, the virtual data system providing access to an aggregated set of data available from a plurality of server data systems respectively operating within a plurality of server computer systems;

30

receiving a client data access transaction from a client computer system, the client data access transaction specifying a data access operation to be performed relative to the virtual data system presented to the client computer system;

processing the client data access transaction in relation to metadata associated
5 with the virtual data system to perform at least one of:

if the data access operation requires access to data available from the plurality of server data systems, then translating, using the metadata associated with the virtual data system, the client data access transaction into a server data access transaction for access to the data available from one of the plurality of
10 server computer systems according to the data access operation specified by the client data access transaction; and

if the data access operation does not require access to data stored by the plurality of server data systems, then accessing metadata associated with the virtual data system to complete processing of the data access operation specified
15 by the client data access transaction.

30. The method of claim 29 wherein:

the virtual data system is a virtual file system and the server data systems are server file systems respectively associated with each of the plurality of server computer
20 systems;

the aggregated set of data of the virtual file system is a collective set of files stored within each of the server file systems of each of the plurality of server computer systems;

the client data access transaction is a file system request for access to information
25 associated with a file presented to the client computer system within the virtual file system; and

wherein the metadata is information maintained by the adaptive load balancer to map virtual file parameters within file system requests provided from client computer systems for access to information associated with a file presented within the virtual file
30 system to corresponding physical file parameters of an instance of a file in at least one of the server file system maintained by at least one of the server computer systems.

31. The method of claim 29 wherein:

the virtual data system is a virtual web service system and the server data systems are web services data systems respectively associated with each of the plurality of server computer systems;

5 the aggregated set of data of the virtual data system is a collective set of service objects stored within each of the web services data systems of each of the plurality of server computer systems;

the client data access transaction is a web request for access to information associated with a service object presented to the client computer system within the virtual
10 web service system; and

wherein the metadata is information maintained by the adaptive load balancer to map virtual object parameters within web requests provided from client computer systems for access to information associated with a service object presented within the virtual data system to corresponding physical object parameters of an instance of a
15 service object in at least one of the web services data systems maintained by at least one of the server computer systems.

32. The method of claim 29 wherein the adaptive load balancer services client data transactions and operates in conjunction with the server file systems of each of the server
20 computer systems according to at least one of the NFS protocol and the CIFS protocol.

33. An adaptive load balancer comprising:

a processor;

a memory;

25 a first communications interface capable of communicating with a client computer system;

a second communications interface capable of communicating with a plurality of server computer systems; and

an interconnection mechanism coupling the processor, the memory, the first
30 communications interface and the second communications interface;

wherein the memory is encoded with an adaptive transaction application that when performed on the processor, produces an adaptive transaction processor that causes the adaptive load balancer to perform the operations of:

presenting a virtual file system to the client computer system via the first
5 communications interface, the virtual file system providing access to an aggregated set of files stored by a plurality of server file systems respectively associated with the plurality of server computer systems;

receiving a client file access transaction from a client computer system via the first communications interface, the client file access transaction specifying a file access
10 operation to be performed relative to the virtual file system presented to the client computer system;

processing the client file access transaction in relation to metadata encoded in the memory associated with the virtual file system to perform at least one of:

if the file access operation requires access to a file stored by the plurality
15 of server file systems, then translating, using the metadata associated with the virtual file system, the client file access transaction into a server file access transaction for access to the file maintained within one of the plurality of server computer systems according to the file access operation specified by the client file access transaction over the second communications interface; and

20 if the file access operation does not require access to a file stored by the plurality of server file systems, then accessing the metadata in the memory associated with the virtual file system to complete processing of the file access operation specified by the client file access transaction.

25 34. The adaptive load balancer of claim 33 wherein operation of the adaptive transaction processor causes the adaptive load balancer to determine if the file access operation requires access to a file stored by the plurality of server file systems, and if so, causes the adaptive load balancer to perform the operations of:

forwarding the server file access transaction over the second communications
30 interface to a selected one of the plurality of server computer systems specified by the metadata for processing of the file access operation specified by the client file access transaction by the selected one of the plurality of server computer systems;

receiving, via the second communications interface, a server transaction response from the server computer system to which the server file access transaction was forwarded;

5 translating, using the metadata associated with the virtual file system in the memory, the server transaction response into a client file access response; and forwarding, over the first communications interface, the client file access response to the client computer system from which the client file access transaction was received.

35. A computer program product having a computer-readable medium including
10 adaptive transaction application computer program logic encoded thereon that, when performed in an adaptive load balancer, operates to process client data access transactions via the operations of:

presenting a virtual data system to a client computer system, the virtual data system providing access to an aggregated set of data available from a plurality of server
15 data systems respectively operating within a plurality of server computer systems; receiving a client data access transaction from a client computer system, the client data access transaction specifying a data access operation to be performed relative to the virtual data system presented to the client computer system; processing the client data access transaction in relation to metadata associated
20 with the virtual data system to perform at least one of:

if the data access operation requires access to data available from the plurality of server data systems, then translating, using the metadata associated with the virtual data system, the client data access transaction into a server data access transaction for access to the data available from one of the plurality of
25 server computer systems according to the data access operation specified by the client data access transaction; and

if the data access operation does not require access to data stored by the plurality of server data systems, then accessing metadata associated with the virtual data system to complete processing of the data access operation specified
30 by the client data access transaction.

36. The computer program product of claim 35 further encoded with logic instructions that cause the adaptive load balancer to determine if the file access operation requires access to a file stored by the plurality of server file systems, and if so, to cause the adaptive load balancer to perform the operations of:

5 forwarding the server file access transaction to a selected one of the plurality of server computer systems specified by the metadata for processing of the file access operation specified by the client file access transaction by the selected one of the plurality of server computer systems;

receiving a server transaction response from the server computer system to which
10 the server file access transaction was forwarded;

translating, using the metadata associated with the virtual file system, the server transaction response into a client file access response; and

forwarding the client file access response to the client computer system from
which the client file access transaction was received.

15